

FAULT-DETECTING COMPUTER SYSTEM

BACKGROUND

[0001] Two classes of hardware-related errors are considered to occur in computational systems: hard errors and soft errors. A hard error is manifested as an improper behavior of the operation of a computer system that persists and continues to cause the system to produce improper behavior and results for a significant period after an initial error occurs. A soft error is a non-recurring error generated by a temporary anomaly in a computer hardware device. Soft errors involve an improper behavior of the computer system that does not persist beyond a certain period of time. After this time has elapsed further operation of the system proceeds normally.

[0002] As the physical devices that make up computer systems have become smaller and more numerous, many recurring physical phenomena are now more likely to cause temporary faults in the operation of these devices resulting in the disruption of the operation of the digital logic and state making up a computing system, often resulting in soft errors. Soft errors are generally more difficult to detect than hard errors. Soft errors are assumed to be more frequent than hard errors and are also assumed to occur sufficiently often that their effect should be considered in computer systems design. Undetected soft errors can result in incorrect results being reported as the result of a computation, corrupt data being stored to disk or other persistent media, or transmitted over network connections, or result in anomalous behavior of a program or of the entire computer system. It is desirable to provide error detection coverage for the subsystems of the computer system architecture which have the highest error rates using techniques which provide detection of soft errors and, optionally, of hard errors. These subsystems typically include the system main memory, the various levels of processor caches as well as system TLB (translation lookaside buffers), I/O and interconnection 'fabric'. When an error is detected it is often desirable to provide a way of correcting the error so that the

computation can continue to produce a correct result. If an error occurs in one of these subsystems, the error will be detected and corrected before it is delivered to other subsystems, thereby obviating the need for the error to be addressed by the other subsystems. This leaves the uncovered subsystems to be addressed. In many computer system designs large portions of the central processing unit are not covered by error detection or error correction.

[0003] With the continuing development of VLSI processors having ever-increasing component density, the susceptibility of these processors to 'soft' errors caused by sources such as cosmic rays and alpha particles is becoming an issue in the design of computational systems. Error detecting and correcting codes are widely applied to the design of computer system memory, caches and interconnection fabric to verify correct operation and to provide correction of the representation of data in the event that either soft or hard errors occur. Protecting the processor electronics is a more difficult task since a processor has many more structures of greater complexity and variety than computer memory devices. Existing hardware techniques for protecting the processor electronics require the design and incorporation of significant logical structures to check, contain and recover from errors which might occur in the core structures that make up the processor.

[0004] Other processor-oriented error detection techniques have included providing multiple processors running the same instructions in 'lock step' and associated self-checking hardware to verify that all results visible externally from each processor match the results of each (or a majority) of its peers to ensure correct operation. In implementation of these techniques where the comparisons do not match, additional complexity is required to limit the propagation of any erroneous state. In addition, special procedures must be performed to either rule the result of the computation as invalid or to recover the state of the computation. All of this adds to the cost and complexity of the system design.

[0005] Software techniques have also been proposed to address errors in computation. Some of these techniques involve fully executing a program multiple times and comparing the results, and then re-executing the

computation until the results match. All of the above techniques multiply the computing resources and time required for a computation to complete. Furthermore, some of these techniques will not detect certain classes of hard errors. Other software fault tolerance techniques assume that a computation will fail in such a way that the computation will stop or 'fail fast', or that errors will be detected by error exception checking logic normally incorporated in processor designs. These techniques often provide inadequate coverage of soft errors.

[0006] From the foregoing, it can be seen that methods for detecting improper operation of computer systems often require extensive hardware and software to support the detection of improper operation, to minimize damage resulting from incorrect results due to improper operation, and also to minimize the number and extent of special actions needed to recover and continue processing in the face of a detected fault. Such systems have often employed doubly or triply redundant hardware and extensive checking and correction logic beyond that required for the basic computation environment itself. Alternative software fault tolerance techniques typically require the adoption of specialized programming techniques which can impact the design of system and applications software, or which require multiple executions of a program and subsequent comparison of the results of two or more program executions.

[0007] The implementation of existing techniques for detecting soft errors, either hardware- or software-based, thus requires significant additional hardware, software, and/or other resources.

SUMMARY

[0008] A system is disclosed for detecting computational errors in a digital processor executing a program. Initially, the program is divided into computation segments, and source code for at least one of the segments is compiled to generate two redundant code sections. Comparison code is also generated for comparing results produced by execution of the two code sections. Each of the code sections is then executed in a different

computational domain to generate respective results. The results of the computation are executed to alter further flow of the program only if the respective results are identical.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] Figure 1 is a diagram showing certain components of an exemplary VLIW processor (prior art);

[0010] Figure 2 is a diagram showing exemplary components and process flow for a temporal replication fault detection system;

[0011] Figure 3 is a diagram showing exemplary components and process flow for a spatial replication fault detection system; and

[0012] Figure 4 is a flowchart illustrating exemplary steps performed during operation of the systems shown in Figures 2 and 3.

DETAILED DESCRIPTION

[0013] Related systems of software techniques for detection of digital processor-related errors are described herein. When combined with existing computer architectures, these systems provide effective fault detection coverage for a processor. The term 'processor' is used in this document to refer to central processing units ('CPU's) as well as digital processors providing other types of functionality. The fault detection techniques described herein may also be used to provide efficient recovery from detected fault conditions. In exemplary embodiments, the techniques may be employed without requiring modifications to the architecture, structure, or source code of applications programs.

[0014] Figure 1 is a block diagram of relevant sections of an exemplary VLIW (Very Long [or Large] Instruction Word) processor 101, such as a Intel Itanium II, that is suitable for use in the present system. VLIW describes an instruction-set philosophy in which a compiler packs a number of basic, non-interdependent operations into the same instruction word. When fetched from cache or memory into the processor, these fixed-length words (instructions) are broken up into several shorter-length instructions which are dispatched to independent functional units (also known as 'execution units'),

where they are executed in parallel. In the processor shown in Figure 1, instructions in instruction cache 110 are queued in instruction queue 109, issued via issue ports 108, and executed via functional units 102–105 using associated registers 106A/106B, described below.

[0015] Processor 101 includes two branch/compare units 102A/102B, two integer units 103A/103B, two load/store units 104A/104B, and two floating point units 105A/105B. Each of the functional units has a corresponding register or register set, which is partitioned into two corresponding but separate parts as indicated by partitions 106A and 106B. The two groups of registers 106A/106B are collectively referred to as a 'register file' 107. The present system is capable of functioning without the parallel branch/compare unit 102B, but the examples shown herein assume that two compare units 102A/B are available on processor 101. The use of partitioned registers allows the detection and repair of errors in register file 107 or paths to/from the register file. The present system includes encoding of different register names into redundant instructions (e.g., load, store, compare) to utilize these partitioned registers.

TEMPORAL REPLICATION

[0016] Soft errors that affect a processor are primarily a result of physical phenomena (e.g., alpha particles and cosmic rays) which are observed to occur randomly but which have some average rate of occurrence and a probability distribution of event durations during which a system behaves incorrectly, or during which the state of the system is altered. Furthermore, the disruptions are generally confined to a single active device or a cluster of physically adjacent devices on a VLSI chip. The observation can be made that the mean time between occurrences of these events is much greater than the maximum duration of disruption. Furthermore, the probability that the same circuit will be disrupted in the same way by a second event after the effects of the first event have ended is also extremely small; as a result, the possibility of two independent identical sequential errors occurring in the same computation units close together in time can be

neglected. Therefore, the technique of temporal replication can be used to create multiple computing domains that can be employed to verify that the computation has not been disrupted in a significant way by a soft error.

[0017] From the probability distribution of event durations, a maximum period of disruption, D_{max} , can be identified such that the probability that an event duration will be longer than D_{max} is small enough that longer durations need not be considered. The average maximum duration of disruptive events due to cosmic rays, alpha particles and other randomly occurring disruptive phenomena dictates a value for D_{max} equal to some predeterminable number of processor clock cycles. The duration of these disruptive events is a function of the particle type and energy along with the properties of the semiconductor processes and design of the devices on the processor chip. Therefore, the actual value for D_{max} may be determined for any real processor design. The applicable value for D_{max} for a particular processor may be determined by detailed simulation of the soft error causes as part of the design processes, determined by measurement of populations of actual devices that make up the processor, or determined through accelerated error rate measurement techniques. For example, for processors with clock frequencies of approximately 1 gigahertz, D_{max} may have a value of several CPU clock cycles.

[0018] Figure 2 is a diagram showing exemplary components and process flow for a temporal replication fault detection system 200. As shown in Figure 2, the source code 201 for a program of interest is separated into computation segments 207 by compiler 202 based on a model wherein each segment takes a set of inputs, performs computations on the input values, and exposes a set of outputs to further computation. Each code segment is processed by compiler 202 and the resultant generated code 203 is passed to an optimizer 205, which schedules the execution of operations in order to best make use of a particular processor's available resources.

[0019] The present method requires no significant modifications to be implemented in a typical compiler prior to the code generation phase. One possible modification comprises the processing of a compiler flag to turn error

checking on or off. In the code generation phase (which follows the source code parsing phase), compiler 202 reads an intermediate encoding of the program semantics and performs standard operations including allocating space and resources on the object computer, such as assigning addresses to identifiers and temporaries. In addition to these operations performed by typical compilers, compiler 202 also generates code for operations that allocate and reallocate resources (such as registers), to hold temporary values.

[0020] The code generation phase of compiler 202 is modified to generate error handling code 204 which verifies the correct operation of each segment of the program as it is executed. The resources of processor 101 are used in such a manner that the redundant and checking computations are each performed in a different computational domain from the domain performing the initial computation. The error handling code 204 generated by compiler 202 is further structured so that an appropriate action for error containment is taken, and, in an alternative embodiment, recovery action is initiated upon detection of an error.

[0021] Present processors typically incorporate multiple execution units in their design to improve processor performance. Multiple, or redundant, execution units are typically present in both multiple issue architectures such as HPPA ('Hewlett-Packard Precision Architecture') or SPARC (Scalar Processor ARChitecture), and also in VLIW architectures such as EPIC IPF ('Explicitly Parallel Instruction Computer Itanium Processor Family'). Frequently, the execution units are not fully utilized due to serialization with I/O and memory operations. As a result, it is often possible to schedule the execution of redundant checking calculations without significant impact on program execution time. Control over the scheduling of these resources is typically not provided in multiple issue architectures and may not be explicit in the case of some VLIW designs; therefore, resource scheduling is performed by optimizer 205. Optimizer 205 reorders the code and schedules the execution of operations in order to best make use of a

processor's available resources, such as functional units, timings and latencies.

[0022] In the present temporal replication method, optimizer 205 schedules execution of redundant code sections 210/215 so that a minimum number of clock cycles (i.e., a minimum amount of time) will elapse between the execution of primary copy 210 and secondary copy 215 of a particular segment of compiled source code. The order of execution of the copies is not important as long as the time between the use of the same hardware resource 206 by the primary/secondary pair of code sections 210/215 is greater than some delta, e.g., Dmax. Given that Dmax is known at the time a program is being compiled to run on a certain processor, compiler 202 in the present system 200 ensures that each section of code 215 that performs the redundant calculations and checking is executed at least Dmax processor cycles apart from the section of code 210 that performs the initial calculation/checking. Optimizer 205 may insert no-ops ('Nops') or schedule other operations between the two sections of code 210/215 to ensure proper spacing of the execution in time.

[0023] In an alternative embodiment, a mechanism is provided to incorporate the length of time corresponding to Dmax in a way that can be interrogated by programs running on processor 101. For example, the value of Dmax may be used by these programs (other than compiler 202), to time skew the execution of redundant threads accordingly to allow for an amount of wait time approximately equal to Dmax.

[0024] The compiled code shown in Table 1 below is an example showing how the operation $A+B=C$ might be performed on an exemplary VLIW processor, such as processor 101. In the example shown in Table 1, the VLIW processor allows a five operations per cycle in its instruction word; only one branch/compare unit is shown. The instructions shown in each row in Table 1 are issued every clock cycle unless the processor stalls waiting for an operand. In the example below, 'BRUnit' is a branch/compare unit (102A/102B), 'ALU/cmpU' is an integer unit (103A/103B), 'Load/storeU' is a load/store unit (104A/104B), and R1–R3 are registers (106A/106B). The

VLIW processor characteristics indicated above are also applicable to all subsequent examples set forth below.

TABLE 1

Cycle	BRU	ALU/cmpU	Load/storeU	ALU/cmpU	Load/storeU
(1)	Nop	Nop	Load R1=A	Nop	Load R2=B
(2)	Nop	Nop	Nop	Nop	Nop
(3)	Nop	Nop	Nop	Add R3=R1+R2	NoP
(4)	Nop	Nop	Nop	Store R3, C	Nop

[0025] Although the above processor is capable of parallelism, there are still a number of Nops in the compiled code shown above. Compiler 202 may include code to schedule instructions in the available slots and issue pre-fetches, etc., in order to increase performance by scheduling more operations per cycle and by reducing latency.

[0026] The compiled and optimized code shown in Table 2 below is an example of the present method of temporal replication for performing the $A+B=C$ operation shown in Table 1. As shown in Table 2, the operation of loading registers R1 and R2 with values of A and B, respectively, is first performed in clock cycle 1, and is repeated at a later time using the same registers in clock cycle 4. The result of the first addition operation is saved in register R3 in cycle 2 and compared by verification code 204, at step 220, against the result of the second addition operation (stored in register R4 in cycle 5). If the values stored in registers R3 and R4 are not equal, a branch to an error handling routine 230 is taken, otherwise, processing continues with the next segment of code at step 225. Compiler 202 breaks the program into segments 207 so that checking of the results of the two operations are checked before the results are 'exposed', or used to alter the further flow of execution of the program.

[0027] Results may be exposed by writing them to an I/O device, or by writing them to a memory area that might be seen by another process or processor or executing a conditional branch which may or may not alter the flow of control in the program. If error recovery is to be implemented, an

additional constraint on a segment 207 is that a segment does not destroy its inputs until checking is successfully completed.

[0028] Optimizer 205 may allow code from adjacent computation segments for executing other program statements to overlap the execution and checking code for a segment such as described in Table 2, provided that the temporal redundancy of each statement is individually maintained and ordering is maintained so that results are exposed in program order and only after the checking sequences for each statement have been successfully executed.

[0029] Error handling routine 230 may provide for retrying an erroneous operation a predetermined number of times, or, alternatively, may cause a fault or take other action in lieu of retrying the operation.

[0030] In the example shown in Table 2, Nops have been inserted into clock cycle 3 by compiler 202. The number of clock cycles that are placed between the execution of the primary copy 210 and the secondary copy 215 of the segment of compiled code is a function of the value of Dmax for a particular processor, as explained above. In the Table 2 example, registers R1–R4 correspond to registers in register file 106A/106B in Figure 1, and 'Error' is the label of error-handling routine 230. The code shown in Table 2 has been compiled/optimized for a Dmax of 3 cycles; that is, there are three clock cycles between the execution of redundant code sections. For example, the first 'Load R1=A' operation has been compiled to execute during clock cycle (1), and the redundant execution of this same operation has been compiled to execute 3 cycles later, during clock cycle (4).

TABLE 2

<u>Cycle</u>	<u>BRUnit</u>	<u>ALU/cmpU</u>	<u>Load/storeU</u>	<u>ALU/cmpU</u>	<u>Load/storeU</u>
(1)	Nop	Nop	Load R1=A	Nop	Load R2=B
(2)	Nop	Add R3=R1+R2	Nop	Nop	Nop
(3)	Nop	Nop	Nop	Nop	Nop
(4)	Nop	Nop	Load R1=A	Nop	Load R2=B
(5)	Nop	Add R4=R1+R2	Nop	Nop	Store R3,C
(6)	Nop	Comp R4,R3	Nop	Nop	Nop

(7) **BNE Error**

Label Error: //Retry and error handling routine

ERROR HANDLING

[0031] The present system performs one or more checks, as indicated by decision block 220 in Figure 2 (and by block 320 in Figure 3), to ensure that the results of computations performed for a code section in two independent computation domains (i.e., temporal or spatial domains) are identical prior to exposing the code to further computation, or before using the result to direct a branch operation. This can be done both before and after the branch is actually taken in order to provide opportunities for optimization by optimizer 205.

[0032] In the case that a mismatch is found between the redundant computations the program will branch to error handling code 230. Recovery may be as simple as indicating an error and terminating the execution of the program ('fail fast'). This technique may be adequate if other levels of recovery are provided by the system. Alternatively, the program may be restarted from the beginning, although this procedure may not be acceptable in some kinds of interactive applications.

[0033] In a more comprehensive recovery procedure, the last program segment is re-executed. Since no computed values are exposed until all computations are checked, a program stage, or segment, that produces an erroneous result may be safely re-executed from the beginning to recover from an error. In an alternative embodiment, a flag is set, indicating that an error recovery operation is in progress. This flag is cleared if the stage of the computation completes successfully. If a second error is encountered in attempting to execute this stage of the program, an indication will be given that a hard error has been encountered.

[0034] A further alternative error handling technique includes structuring a program so that the results are computed three or more times on different domains, wherein the program code is structured so that the computed results delivered by the majority of the computational domains is

exposed as the result of execution of a particular segment of the code. Note that any of these methods may be selectively used on only the code needing this level of protection.

SPATIAL REPLICATION

[0035] Figure 3 is a diagram showing exemplary components and process flow for a spatial replication fault detection system 300. In the spatial replication method, the code for a particular program may be executed two or more times, with each execution path using separate processor resources including functional units and registers.

[0036] As shown in Figure 3, source code 201 for the program of interest is initially separated into computation segments 207 by compiler 202 based on a model wherein each segment takes a set of inputs, performs computations on the input values, and exposes a set of outputs to further computation. Each segment 207 is then compiled by compiler 202 to generate primary compiled code section 303, secondary compiled code section 305, and verification code 304, which are input to optimizer 205.

[0037] Optimizer 205 then generates primary and secondary copies 306 and 307 of code corresponding to code segments 303 and 305, ensuring that the primary and secondary copies are executed using different registers and functional units 206(1)/206(2) (as well as any other execution resources utilized by the program), respectively. An optimized copy 308 of verification code 304 is also generated by optimizer 205. Optimizer 205 may reorder the operations in copies 306/307 as long as it does not eliminate the primary or secondary operation in a pair, or cause them to be executed on the same resource or expose a result until the checking code has verified that the results are identical. As in the temporal replication case, the compiler divides code into segments so that checking is performed prior to the exposure of any results or transfers of control.

[0038] In an alternative embodiment, 'hints', such as tags, are provided to allow a computation domain to be associated with a set of execution and data resources that is disjoint from the resources used for other

computation domains. Execution of different computation domains may be allowed to overlap. Compiler 202 generates a first code segment for the execution of each stage of program execution along with a second code segment for redundant computation and checking of the results, specifying a different computation domain for each segment.

[0039] For example, if in a stage of the execution of a program one needed to assign $s = \sin(a)$, compiler 202 may generate code such as the following:

```
ts=sin(a); in the first computation domain;  
rs=sin(a); in the second computation domain; and  
if (ts<>rs) go to recovery routine;  
[at this point s can be exposed for further computation]  
s=ts;
```

[0040] Code is then generated for subsequent program stages. In the present embodiment, one computation domain uses a first set of execution units, e.g., functional units 102A–105A and registers in register group 106A, and the other computation domain uses a second set of execution units, e.g., functional units 102B–105B and registers in register group 106B.

[0041] 'Hints' maybe given to the processor's execution unit scheduling logic by tagging groups of instructions with information that indicates the computation domain with which each group of instructions is associated. Hints may also be included in the compiled code to assure that the optimizer or hardware does not discard the duplicated operations.

[0042] Before the result of a computation is exposed, or used for further computation, checking of the redundant results is performed. Verification code 304 generated by compiler 202 is executed, as indicated by decision block 310, to compare the results of execution of primary and secondary code copies 306/307. This checking may also be performed in a computation domain different from those used in the actual computation. In the case that the results do not match, recovery actions can be attempted that

are similar to those described with respect to Figure 2 in the above section on temporal replication. Compiler 202 can make use of the explicit scheduling available in the instruction set of many VLIW processors to ensure that redundant pairs of code are not executed by the same functional units. If a discrepancy in results is found, appropriate recovery action is taken by error handling routine 320. This recovery action may include re-execution, failing, or trapping to software or operating system handlers.

[0043] The compiled code shown in Table 3 below is an example of the present method of spatial replication for performing the $A+B=C$ operation shown in Table 1. As shown in Table 3, registers R1 and R11 are loaded with the value of A in clock cycle 1, and registers R2 and R12 are loaded with the value of B in clock cycle 2. Registers R1 and R2 are, for example, part of register group 106A and registers R11 and R12 are part of register group 106B. During clock cycle 3, registers R3 and R13 are used to sum the contents of registers R1/R2 and R11/R12, respectively.

[0044] Register R4 is then loaded with the stored value of 'C', and the result of the first addition operation is then compared by verification copy 304/308 in clock cycle 4 (step 310 in Figure 3), against the result of the second addition operation. If the values stored in registers R3 and R13 are not equal, a branch to error handling routine 320 is taken, in cycle 5. During clock cycle 6, the sum stored in register R13 is stored in processor memory as 'C'. If the values stored in registers R3 and R13 match, then the values stored in registers 3 and 4 are compared, in clock cycle 7. Here, the value of an operand stored to memory is reloaded and its fetched value compared to that which was supposed to be stored. This is done to be sure that there was no error on the paths from the register to memory or in the memory controller. If the values stored in registers R3 and R4 are not equal, a branch to error handling routine 320 is taken in cycle 8, otherwise, processing continues with the next segment of code, at step 315.

[0045] The results of the two operations are checked before the results are 'exposed', or used to alter the further flow of execution of the program. Error handling routine 320 may provide for any combination of the

following actions: retrying an erroneous operation a predetermined number of times; causing a fault, or taking other action in lieu of retrying the operation; error reporting; and statistics collection.

[0046] Each column of instructions in Table 3 is executed by a specific functional unit in processor 101, as well as by a specific group of registers, in either group 106A or 106B in the register file 107. Register file 107 is partitioned such that the same register resources are not used by the primary and secondary code copies 306/307.

TABLE 3

<u>Cycle</u>	<u>BRUnit</u>	<u>Primary units</u>		<u>Secondary Units</u>	
		<u>ALU/cmpU</u>	<u>Load/storeU</u>	<u>ALU/cmpU</u>	<u>Load/storeU</u>
(1)	Nop	Nop	Load R1=A	Nop	Load R11=A
(2)	Nop	Nop	Load R2=B	Nop	Load R12=B
(3)	Nop	Add R3=R1+R2	Nop	Add R13=R11+R12	NoP
(4)	Nop	Nop	Nop	Comp R13,R3	Nop
(5)	BNE Error	Nop	Nop	Store R13, C	Nop
(6)	Nop	Nop	Load R4=C	Nop	Nop
(7)	Nop	Nop	Comp R4,R3	Nop	Nop
(8)	BNE Error	Nop	Nop	Nop	Nop

Label Error: // Error and retry handling routine

[0047] Note that optimizer 205 may schedule subsequent operations into some of the Nop spots in the code shown above. As shown in the example in Table 3, duplicated code using different result registers allows comparison of results to determine if there was an error in the functional units, registers, or on the paths between them. The same is true of compare operations as well.

[0048] In an alternative embodiment, the target address or label of a branch (or other change of control operation) may be loaded into a register so that a determination can be made as to whether the change of control was correctly executed, by comparing the value stored in the register with a literal value of the address associated with the label to which the branch was taken.

The value stored and compared need not be the address but a value that is sufficiently unique to the label or entry point such that it is unlikely that an errant branch would take control with an identically encoded label or entry point.

[0049] In a further alternative embodiment, parameters in procedure and system calls may be duplicated, including passing a redundant return address or command code. Similarly, duplicated results may be returned. These techniques help ensure that the parameters to, and results from, a called routine are correct.

[0050] If the host system does not have adequate error detection and correction for memory and the paths to and from memory, two separate data regions, as represented by primary and secondary code copies 306/307, may also be maintained. Data is fetched from the redundant areas and compared to assure the fidelity of the data.

[0051] In an alternative embodiment, rather than comparing the results of two spatially distinct computations and branching to an error handling routine, or re-executing the code, the code for a particular program may be executed in more than two spatial domains and the results voted on to determine which result (i.e., the majority result, or consensus) is to be executed.

[0052] Figure 4 is a flowchart illustrating exemplary steps performed during operation of the systems shown in Figures 2 and 3. As shown in Figure 4, in step 405, source code for a program is first segmented into computation segments, at step 406, and then compiled and optimized in one of two forms. In either form, the resulting compiled object code will perform a redundant computation in a different computational domain from the domain performing the initial computation.

[0053] If the resultant compiled code is to be executed in a time-skewed manner (as described above with respect to Figure 2), then at step 407, compiler 202 and optimizer 205 generate and schedule execution of two redundant code segments so that a minimum number of clock cycles will

elapse between the execution of primary copy 210 and secondary copy 215 of a particular segment of the compiled source code.

[0054] If the compiled code is to be executed via different hardware entities, then at step 408, compiler 202/optimizer 205 generate essentially redundant primary and secondary copies 306 and 307 of a particular segment of code, ensuring that the primary and secondary copies use different registers and functional units 206(1)/206(2). These two copies are said to be essentially redundant because, although the two copies are functionally identical and perform the same computation(s), the two copies are not strictly identical, since different registers and functional units are employed in the execution of each copy. It is to be noted that In either of the above cases (described in steps 407 and 408), compiler 202 may be configured to perform the additional functions of optimizer 205, as described herein.

[0055] Verification code is generated by compiler 202 at step 410, during compilation of the corresponding code segment. At step 420, the redundant copies of a compiled code segment are executed by processor 101. The verification code generated in step 407 is executed at step 425 to compare the respective results of execution of primary and secondary copies 306/307. At decision block 430, if a discrepancy in results is found, appropriate action is taken by the appropriate error handling routine 230/320. This error recovery action may include re-execution (N1 – step 433), or failing or trapping (N2 – step 432) to software or operating system handlers. If the respective results of execution of primary and secondary copies 306/307 are identical, then at step 434 the results are committed, and redundant copies of the next segment of code are executed, at step 420.

[0056] In an alternative embodiment, the verification code itself generated by the compiler may be constructed so that verification is executed redundantly in multiple computation domains.

[0057] The above-described operations can be implemented in a standard compiler, or in a tool that dynamically translates code to native machine code or object format such as is done in 'just in time' (JIT) compilers. In another implementation or tool, software that performs static or dynamic

code reorganization or optimization may be employed to dynamically translate legacy code into a redundant form, or incrementally translate existing code, in accordance with the present method. A design compliant with the present system may use all or some of the techniques above as determined by the amount of protection that is desired, as well as the performance requirements of the code, and as appropriate to augment whatever error detection mechanisms are built into the relevant hardware.

[0058] Instructions that perform the operations described with respect to Figures 2–4 may be stored on computer-readable storage media. These instructions may be retrieved and executed by a processor, such as processor 101 of Figure 1, to direct the processor to operate in accordance with the present system. The instructions may also be stored in firmware. Examples of storage media include memory devices, tapes, disks, integrated circuits, and servers.

[0059] Certain changes may be made in the above methods and systems without departing from the scope of the present system. It is to be noted that all matter contained in the above description or shown in the accompanying drawings is to be interpreted as illustrative and not in a limiting sense. For example, the processor shown in Figure 1 may be constructed to include components other than those shown therein, and the components may be arranged in other configurations. The elements and steps shown in Figures 2–4 may also be modified in accordance with the methods described herein, without departing from the spirit of the system thus described.